

Understanding and philosophical methodology

Magdalena Balcerak Jackson ·
Brendan Balcerak Jackson

© Springer Science+Business Media B.V. 2011

Abstract According to Conceptualism, philosophy is an independent discipline that can be pursued from the armchair because philosophy seeks truths that can be discovered purely on the basis of our understanding of expressions and the concepts they express. In his recent book, *The Philosophy of Philosophy*, Timothy Williamson argues that while philosophy can indeed be pursued from the armchair, we should reject any form of Conceptualism. In this paper, we show that Williamson’s arguments against Conceptualism are not successful, and we sketch a way to understand understanding that shows that there is a clear sense in which we can indeed come to know the answers to (many) philosophical questions purely on the basis of understanding.

Keywords Understanding · Metaphilosophy · Williamson, Timothy · Conceptual truth · Conceptual analysis · Analyticity

1 Introduction

Throughout most of its history, philosophy has been thought of and conducted as a discipline that is importantly distinct from the empirical sciences, and as a discipline that can be pursued “from the armchair.” What justifies this attitude towards philosophy? Philosophers often adopt it because they are what we will call *Conceptualists*, philosophers who, as Timothy Williamson puts it, believe that “whatever can be achieved through exercise of [linguistic/conceptual] competence

M. Balcerak Jackson · B. Balcerak Jackson (✉)
Emmy Noether Research Group, Philosophy Department, University of Cologne,
Richard-Strauss-Strasse 2, 50931 Cologne, Germany
e-mail: bbalcerakjackson@gmail.com

M. Balcerak Jackson
e-mail: mbalcerakjackson@me.com

and reflection thereon will be a feasible goal for philosophy.”¹ But while Williamson endorses the conception of philosophy as an armchair discipline that Conceptualism is often used to justify, he thinks Conceptualism itself is a mistake. One of the main purposes of his recent book, *The Philosophy of Philosophy*, is to argue that Conceptualism in any form should be rejected.

Conceptualism, as we think of it, is an attempt to answer the question: How can philosophers make progress on, say, questions about knowledge or free will without going out and doing empirical investigation, but rather just by engaging in armchair reflection? According to the Conceptualist answer, the philosopher can operate from the armchair because she is trying to uncover whatever can be learned about these questions merely on the basis of her understanding of expressions like ‘knowledge’ and ‘free will’ and of the concepts they express.² But anyone who endorses Conceptualism has a further question to answer: how can we come to know the answers to philosophical questions just on the basis of understanding? Williamson’s critique of Conceptualism is an inductive argument that proceeds by attacking various attempts to provide an answer to this question. He concludes from this that Conceptualism probably cannot adequately explain the legitimacy of armchair philosophy, and that we ought to look elsewhere for such an explanation.

We disagree with this conclusion. In our view, Conceptualism has the resources to provide a good justification (or at least a central part of such a justification) for armchair philosophy. Our goal in this essay is to evaluate Williamson’s argument and explain where we think it goes wrong.

Our discussion will proceed as follows. Williamson considers three attempts to explain how we can come to know the answers to philosophical questions on the basis of understanding, and he argues that each of the three attempts faces fatal difficulties.³ The first two attempts approach the question indirectly: they try to identify a special, distinctive class of truths to which we plausibly have access on the basis of understanding alone, and then they propose that the goal of philosophical inquiry is to uncover truths of that sort. Williamson raises objections to both attempts to characterize the relevant class of truths. In Sect. 2 we examine those objections and argue that they pose no serious threat to Conceptualism. One reason is that the objections only challenge some old-fashioned (and in our view misguided) ways of articulating Conceptualism that contemporary Conceptualists no longer endorse. But a deeper reason is that, at best, they suggest only that the search for a special class of truths is a mistake; they do not provide any reason for thinking that armchair philosophical knowledge cannot be grounded in understanding. The closest Williamson comes to giving any such reason is in his argument against the third attempt to justify Conceptualism. This attempt rests on the claim that acceptance of certain truths is necessary for understanding one or more of the concepts contained in them, and Williamson argues that there are no such truths.

¹ Williamson (2007, p. 2). (All page references hereafter will be to this book unless otherwise indicated.).

² Throughout his book Williamson formulates claims both at the level of language and at the level of thought. For ease of discussion we will allow ourselves to move back and forth between the two, without being very careful about the difference when it doesn’t matter.

³ The main negative part of *The Philosophy of Philosophy* (Chaps. 1–4) comprises the discussion of these three attempts.

However, in Sect. 3 we argue that this argument does not show that philosophical knowledge cannot be grounded in understanding. It is quite plausible that the understanding of expressions and concepts is (or involves) a cognitive capacity, and the reflective exercise of this capacity can provide the basis for coming to know philosophical truths quite independently of whether or not there are any beliefs that are necessary for understanding. Finally, for reasons that will become clear in Sect. 3, we suspect that Williamson is at least largely motivated by a certain conception of understanding, one that is deeply at odds with any sort of Conceptualist approach to justifying armchair philosophy. Therefore in Sect. 4 we try to spell out Williamson's conception of understanding and explain why we find it unconvincing.

Three brief points of clarification are needed before we begin. First, as we will see in Sects. 2 and 3, much of Williamson's discussion in *The Philosophy of Philosophy* is framed in terms of the distinction between metaphysical and epistemological conceptions of analyticity first explicitly drawn by Paul Boghossian as part of a project to rehabilitate the notion of analyticity.⁴ Boghossian argues that the epistemological conception of analyticity is not susceptible to the worries that have plagued the metaphysical conception, and that have led many philosophers to abandon analyticity altogether. Williamson's discussion in *The Philosophy of Philosophy* is, in part, a response to Boghossian's project. The core of Williamson's argument against Conceptualism—which is the focus of Sect. 3 here—is an argument against the existence of truths that are epistemologically analytic in (something closely related to) Boghossian's sense. Since we maintain that Williamson's argument fails, the question arises of how our defense of Conceptualism relates to Boghossian's project. Our main aim here is to show that Williamson does not succeed in ruling out the possibility of acquiring knowledge solely on the basis of understanding. To the extent that our discussion makes it plausible that one *can* ground knowledge in understanding, it contributes to Boghossian's project of developing and defending an epistemological conception of analyticity. However, as we will see, our defense requires thinking in very different ways than Boghossian and other Conceptualists have so far about the link between understanding, belief and knowledge.

Second, in addition to arguing against Conceptualism, Williamson also hopes to show that even though philosophy is legitimately done from the armchair, its results are not generally justified a priori. In fact, for Williamson, “even the distinction between the a priori and the a posteriori turns out to obscure underlying similarities” when we try to apply it to philosophical armchair knowledge.⁵ Conceptualism—and hence our defense of it—is clearly relevant to the question of whether armchair philosophical knowledge is a priori. But here we will not take up that question directly.

Third, in what we might call its “pure” form, Conceptualism is intended as a *complete* account of the epistemology and methodology of philosophy. The pure Conceptualist regards the reflective exercise of one's understanding of expressions

⁴ Boghossian (1996).

⁵ p. 2.

and concepts as the only legitimate way to acquire genuinely philosophical knowledge. We do not wish to defend pure Conceptualism. We have no desire to say, for example, that the philosopher of biology working on levels of selection in evolutionary theory is not doing genuine philosophy (or else that her work should be restricted to analyzing biological concepts). We also want to leave it open whether there are other legitimate armchair methods in philosophy that are not understanding-based. For example, we do not wish to dismiss as illegitimate the phenomenologist's investigation of the structure of experience. Nevertheless, in what follows we will usually frame our discussion in pure Conceptualist terms, merely for ease of expression. What Williamson denies, and what we wish to defend, is that understanding can provide an epistemic basis for arriving at answers to many central philosophical questions.

2 Williamson's argument, part 1

In this section we take up Williamson's arguments against the first two of the three Conceptualist theses that are the main focus of the critical part of *The Philosophy of Philosophy*. The first thesis Williamson considers is a quite literal-minded interpretation of the suggestion that conceptual analysis is the investigation of concepts, or of the meanings of expressions:

CP1 Philosophy aims to discover metalinguistic truths about expressions, and/or metaconceptual truths about concepts.

CP1 construes conceptual truths as truths *about* concepts or expressions. The second thesis construes them rather as *metaphysically analytic* truths—where a metaphysically analytic truth is, roughly, one that is true purely in virtue of meaning:

CP2 Philosophy aims to discover metaphysically analytic truths.

We will discuss Williamson's objections to CP1 and CP2 in turn.

According to CP1, the goal of philosophical inquiry is to discover true meta-linguistic claims—for example, claims about the meanings of (or concepts expressed by) expressions like 'knowledge' or 'freedom'. But Williamson argues that it is rarely the case that the questions philosophers are interested in are—either explicitly or implicitly—meta-linguistic questions. He uses the philosophical investigation of vagueness as a case study.⁶ One way to begin such an investigation is to ask, for example, whether Mars was always either dry or not dry. This is *explicitly* a question about Mars and dryness, not a question about the meaning of 'Mars' or 'dry' or any other expression. Nor, according to Williamson, is there any good sense in which the question is *implicitly* a meta-linguistic one. It is true that philosophers' attempts to answer questions like this often lead them to investigate the meanings of predicates like 'dry', and of the logical constants 'every' and 'or'. But this does not provide any useful sense in which the original question is implicitly meta-linguistic. In an analogous way, physicists' attempts to answer

⁶ Chapter 2.

questions about the behavior of particles sometimes lead them to investigate the workings of their measuring instruments, but this does not show that the physicists' original questions were implicitly about their instruments. What holds for research on vagueness holds for other areas of philosophy as well: philosophers engaged in the armchair investigation of questions such as whether knowledge requires belief, or whether freedom is compatible with determinism, are interested in finding out about knowledge or freedom itself, and not—even implicitly—about 'knowledge' or 'freedom'. So Williamson concludes that CP1 is false.⁷

Williamson is right to reject CP1, but this should not be taken as much of a point in favor of his general critique of Conceptualism. No contemporary Conceptualist of whom we are aware would accept CP1, and for good reason. It is true that philosophers often claim to be investigating our words or concepts, but it is very implausible to construe this as a commitment to CP1. It is far more plausible that such philosophers are primarily interested in knowledge or free will itself, just as Williamson insists, but that they think that we can learn something about those things *by* investigating our words or concepts. Frank Jackson—one of Williamson's primary targets in *The Philosophy of Philosophy*—is explicit about this point:

Although metaphysics is about what the world is like, the *questions* we ask when we do metaphysics are framed in a language, and thus we need to attend to what the users of the language mean by the words they employ to ask their questions. ...metaphysicians will not get very far with questions like: Are there *Ks*? Are *Ks* nothing over and above *Js*? and, Is the *K* way the world is fully determined by the *J* way the world is? in the absence of some conception of what counts as a *K*, and what counts as a *J*.⁸

In fact, as we have just seen, Williamson *agrees* that investigating words and concepts can help us answer questions in philosophy, just as investigating instruments can help answer questions in physics. Certainly, Conceptualists have questions to answer about how the investigation of concepts can tell us something about knowledge or free will, and about the role understanding plays in this investigation. We will have more to say about these questions in Sect. 3.⁹ But the important point here is that Conceptualism does not rest on the simple failure to

⁷ While we agree with Williamson's point here, vagueness is an unfortunate choice of philosophical subjects to use as a counterexample to CP1. The vast majority of the research on vagueness focuses on vague expressions (typically predicates), and one of the main goals of this research is to develop a semantic theory that will adequately capture the behavior of vague expressions while avoiding paradox. Such a theory, obviously, is a collection of meta-linguistic claims. Moreover, it is implausible to suggest that the main focus of research on vagueness is really on non-meta-linguistic questions like whether Mars was always either dry or not dry. We doubt that many theorists have any real interest in this question for its own sake. Some theories imply an answer to this question, and vagueness theorists are interested in whether or not this is a good consequence. But for theories that don't imply an answer, the answer depends on empirical facts about water on Mars, and nobody seems to think that it is the business of philosophers to discover such facts.

⁸ Jackson (1998, pp. 30–31).

⁹ Of course, Jackson and other Conceptualists go to great lengths to provide their own answers to these questions.

distinguish knowledge from ‘knowledge’ that Williamson points out in his critique of CP1.¹⁰

Turning to CP2, Williamson’s argument here focuses on the notion of metaphysical analyticity, which Williamson takes up from the work of Paul Boghossian.¹¹ Metaphysically analytic truths are true solely in virtue of their meaning, and not at all in virtue of how things are in the extra-linguistic world—their truth imposes “no genuine constraint on the world.”¹² According to CP2, the goal of philosophy is to discover such truths. On this picture, philosophy can be done from the armchair because there is no need to go out and test metaphysically analytic claims against the world; their truth makes no demands on what the world must be like. The challenge for this picture is to make the notion of metaphysical analyticity coherent, and to do so in a way that supports the intended justification for doing philosophy from the armchair. Williamson considers several attempts to spell out the notion of metaphysical analyticity, and argues that none succeeds in meeting both of these challenges.¹³ Hence he rejects CP2.

We think Williamson is right to reject CP2, but again, this does not cast much doubt on Conceptualism. Williamson is hardly the only philosopher to doubt the coherence of the metaphysical conception of analyticity. Few philosophers since Quine have wanted to defend it.¹⁴ Metaphysical analyticity makes no appearance at all in the work of contemporary Conceptualists such as Jackson and David Chalmers.¹⁵ And contemporary Conceptualists who discuss analyticity, such as Paul Boghossian and Christopher Peacocke, explicitly reject the metaphysical conception.¹⁶ According to Peacocke, for example, “contemporary theorists of the a priori should not be involved with the uninstantiated and uninstantiable notion of ‘true purely in virtue of meaning’.”¹⁷ Boghossian’s purpose in distinguishing the metaphysical conception is to argue that it should be set aside in order to develop a plausible account of both analyticity and Conceptualism about the a priori.¹⁸ We can’t see any reason why Conceptualists need to invoke any notion of metaphysical analyticity nor, consequently, why they should be concerned with Williamson’s objection to CP2.

¹⁰ Goldman (2007) holds a view according to which philosophical intuitions are primarily evidence for claims about the characteristics of our individual mental representations. However, for precisely this reason Goldman thinks that intuitions have very limited and indirect evidential value for philosophy. Hence it would be implausible to construe him as endorsing CP1, or Conceptualism more generally.

¹¹ Boghossian (1996).

¹² p. 53.

¹³ Chapter 3.

¹⁴ Russell (2008) is a rare exception.

¹⁵ Jackson (1998), Chalmers and Jackson (2001).

¹⁶ Boghossian (1996, 2003) and Peacocke (2004). Boghossian says little about matters of philosophical methodology, and so it is not entirely clear whether to group him among the Conceptualists. Nevertheless, he does develop an account of a significant range of a priori knowledge in terms of understanding, and he is clearly one of Williamson’s targets.

¹⁷ Peacocke (2004, p. 27). Peacocke goes on to refer to those who still believe in metaphysical analyticity as a “possibly non-existent class of theorists” (p. 27, n. 28).

¹⁸ See, e.g. Boghossian (1994, 1996).

It must be acknowledged that metaphysical analyticity did play a central role in the views of the logical empiricists, and of A.J. Ayer—one of Williamson’s predecessors to the Wykeham Chair of Logic at Oxford—in particular. In Chapter 1 of *The Philosophy of Philosophy*, Williamson identifies Ayer as an early proponent of Conceptualism and as a target of his critique. The problems Williamson raises for CP2 certainly are problems for Ayer. In fact, Williamson quotes a passage in which Ayer endorses a view that would seem to leave him open to the objections Williamson raises against *both CP2 and CP1*:

The philosopher, as an analyst, is not directly concerned with the physical properties of things. He is concerned only with the way in which we speak about them. In other words, the propositions of philosophy are not factual, but linguistic in character – that is, they do not describe the behavior of physical, or event mental, objects; they express definitions, or the formal consequences of definitions.¹⁹

So in critiquing CP1 and CP2, Williamson is not tilting at windmills: his objections do raise genuine problems for a well-known and historically influential version of Conceptualism.

Be that as it may, Williamson’s arguments against the Ayer-style logical empiricist version of Conceptualism are hardly grounds for Conceptualists these days to rethink their views. The objections exploit the fact that Ayer attempts to justify Conceptualism by identifying a special class of truths—truths that are somehow insubstantial or merely linguistic, rather than being about the extra-linguistic world—that might plausibly be available to us merely on the basis of understanding. Williamson’s objection, in effect, is that it is either implausible that philosophers are after such truths or else unintelligible what those truths are supposed to be. But raising doubts about finding such a special class of truth does not yet threaten the basic guiding picture with which we began: the picture according to which our grasp of concepts or of the meanings of expressions can provide us with epistemic access to philosophical truths from the armchair.

Actually, a brief look at Ayer’s first-order philosophizing, rather than his meta-philosophical pronouncements, can help make this point clear. When arguing for his version of compatibilism about free will, for example, Ayer says:

Let it be granted, then, when we speak of reconciling freedom with determination we are using the word ‘freedom’ in an ordinary sense. It still remains for us to make this usage clear: and perhaps the best way to make it clear is to show what it is that freedom contrasts with.²⁰

Ayer goes on to present his readers with a series of hypothetical cases involving various sorts of free or constrained actions. It is clear that Ayer takes this process to be a way of making use of his and his readers’ grasp of the ordinary sense of ‘freedom’ and other terms. Ayer also clearly expects the results to give us good grounds for accepting his conclusions about free will.

¹⁹ Ayer (1946, p. 57).

²⁰ Ayer (1954, p. 114).

Analogous observations apply to the work of other early twentieth century empiricists, such as the work on confirmation by Carl Hempel and Nelson Goodman. Hempel's Raven Paradox, for example, is something he thinks arises just by employing our ordinary understanding of 'confirmation', and reflection on the paradox is meant to help us discover general principles concerning when some evidence confirms (or disconfirms) a hypothesis.²¹ Goodman's "new riddle" concerning 'grue'-like cases is also generated by reflection on our ordinary understanding of 'confirmation', in this case for the purpose of providing support for his claim about the relationship between confirmation and the lawlikeness of hypotheses.²²

In all these cases we see the basic Conceptualist picture at work, according to which we use our understanding to gain epistemic access to the answers to philosophical questions. In their explicit methodological theorizing, logical empiricists such as Ayer and Hempel claim that the truths we have access to on the basis of understanding must be truths that are (in some sense) merely meta-linguistic, or truths that (in some sense) impose no constraint on the world; these assumptions lead to CP1 and CP2. Williamson's objections may show that the appeals to meta-linguistic truth and metaphysical analyticity are misguided, but the basic Conceptualist picture outlives these positivist assumptions. To challenge the idea that understanding provides us with epistemic access to the things we want to find out about when we do philosophy, Williamson needs to challenge the epistemic powers of understanding directly.

3 Williamson's argument, part 2

The closest Williamson comes to doing so is in Chapter 4 of *The Philosophy of Philosophy*. Here he considers the suggestion that conceptual truth should be construed in terms of *epistemological analyticity* rather than in terms of metaphysical analyticity, a suggestion that closely follows Boghossian's proposal for rehabilitating the analytic/synthetic distinction. Williamson defines epistemological analyticity as follows:

EA A thought *t* is epistemologically analytic if and only if necessarily, anyone who understands *t* assents to *t*.²³

Following the pattern established by CP1 and CP2 yields:

²¹ Hempel (1945).

²² Goodman (1954).

²³ EA is distinct from the notion of epistemological analyticity formulated in Boghossian (1996), although how exactly it differs depends on how one interprets Boghossian's formulations. Boghossian sometimes characterizes a statement as epistemologically analytic when understanding it suffices for having a *justified belief* in its truth, in which case EA is a weaker notion that requires only belief. However, Boghossian sometimes characterizes a statement as epistemologically analytic when understanding it suffices for *being justified* in believing it. This is naturally construed as a claim about propositional justification, about it being epistemically appropriate for the subject to have the belief, whether or not she actually has it.

CP3 Philosophy aims to discover epistemologically analytic truths.

However, there is an important difference between CP1 and CP2 on the one hand, and CP3 on the other. As we remarked in the previous section, CP1 and CP2 attempt to justify Conceptualism by finding truths with a special sort of subject-matter that plausibly makes them amenable to discovery purely on the basis of understanding. But epistemologically analytic truths need not have any special subject-matter; what distinguishes them is not their subject-matter but their relationship to understanding. Let us introduce a further bit of terminology: when the right-hand side of the biconditional in EA holds for a given thought *t*, so that *t* is epistemologically analytic, we'll say that *understanding guarantees assent* for *t*, or equivalently that there is an *understanding-assent guarantee* for *t*.²⁴ So, for example, the thought *that every bachelor is male* is epistemologically analytic just in case understanding the thought guarantees assent for it; there must be no possible subject who grasps the thought but withholds assents from it.²⁵ According to CP3, then, the goal of philosophical inquiry is to uncover truths such that merely understanding them guarantees that one will assent to them. And the problem for the view, according to Williamson, is that there simply aren't any thoughts for which understanding guarantees assent.

How does Williamson hope to establish this? His strategy is to consider several seemingly paradigm examples of truths for which an understanding-assent guarantee ought to obtain, and to argue that in each case there are hypothetical subjects who understand them without assenting to them; if this is right then we should at least begin to have serious doubts that there are any secure understanding-assent guarantees to be found.²⁶

One such truth is the thought *that every vixen is a vixen*. Even though this is an utterly trivial logical truth, Williamson argues that there are possible subjects who violate the understanding-assent guarantee for it. One such subject is Peter, who holds the view that a thought of the form *that every F is a G* can only be true if there is at least one F. In particular, according to Peter the thought *that every vixen is a vixen* is false unless there is at least one vixen. But Peter doesn't believe in vixens; he thinks that the widespread belief in such creatures is the result of government propaganda. Hence he rejects the thought *that every vixen is a vixen*, even though he understands it.²⁷ Another such subject is Stephen. He believes in vixens, but he also believes there are borderline cases of vixenhood, and he has a particular theory of vagueness. His theory implies that a universally generalized thought of the form

²⁴ Williamson talks about understanding-assent "links" rather than guarantees. However, we wish to avoid the misleading suggestion that the strong necessary connection that Williamson criticizes is the only epistemologically interesting link between understanding and assent, or the only link that might be available to the Conceptualist.

²⁵ Following Williamson (pp. 74–75) the notion of assent employed in the characterization of both epistemological analyticity and of understanding-assent guarantees is meant to be dispositional; a subject need not be actively assenting to a thought throughout the time that she grasps it. It should also be understood so that a subject can count as assenting to the thought *that foxes have fur*, for example, even if she does not make the explicitly meta-level judgment that the thought *that foxes have fur* is true.

²⁶ Chapter 4.

²⁷ pp. 86–87.

that every F is a G is true only if there are no borderline cases of F (or G). Knowing that his theory has this consequence, and believing that there are borderline cases of vixenhood, Stephen concludes that the thought *that every vixen is a vixen* is not true and so withholds assent, even though he understands it.²⁸ If either Peter or Stephen is possible, then mere understanding is not sufficient to guarantee that one will assent to the thought *that every vixen is a vixen*.

We will examine Williamson's reasons for thinking that Peter and Stephen really do qualify as understanding the thought *that every vixen is a vixen* in the next section. In this section we want to grant that Peter and Stephen are genuine counterexamples, and consider what follows if there are no understanding-assent guarantees. Does it follow that we cannot come to know anything purely on the basis of understanding? Williamson certainly thinks so. Before introducing Peter and Stephen, he foreshadows the point of the discussion as follows:

In what follows, we will consider more rigorously what is epistemically available simply on the basis of linguistic and conceptual competence. To a first approximation, the answer is: nothing.²⁹

But this is a mistake. Williamson's argument fails to show that understanding cannot be the epistemic basis for armchair philosophical knowledge, as we shall now argue.

One can begin to see the problem for Williamson by noting that no reasonable Conceptualist should endorse CP3, because CP3 would threaten to make philosophy an entirely trivial enterprise. According to CP3, philosophy aims to reveal truths for which understanding guarantees assent. Suppose, then, that we have a candidate claim C . Let us consider whether it is possible to have a substantive—i.e. not merely verbal—debate about whether C is true. If only one of the parties to an apparent debate over whether C is true correctly understands C then the debate is merely verbal—at least, it is if correcting the misunderstanding would cause the disagreement to evaporate. So if a substantive debate over C is possible, it must be possible for there to be a subject who understands C and assents to it, and it must also be possible for there to be a subject who understands C and yet withholds assent from it. But this is just to say that understanding does not *guarantee* assent for C . In other words, any claim over which there can be substantive, non-merely-verbal disagreement is one that is not epistemologically analytic. According to CP3, then, philosophers are not in the business of advancing claims about which there can be substantive, non-verbal debate (at least insofar as they are doing conceptual analysis). Whenever a philosophical thesis is met with disagreement—especially sustained, reflective disagreement—its author should either dismiss her critics as failing to understand her, or else concede that she has gone beyond the bounds of what philosophical methods can deliver. This is not a picture any sensible Conceptualist should want to endorse.³⁰

²⁸ pp. 87–88.

²⁹ p. 77.

³⁰ Williamson's argument against CP3 actually has its origins in an earlier paper (2003) that responds to Boghossian 2003. In this paper Boghossian appeals to understanding-assent guarantees to help account

The trouble is that Williamson's objection assumes, implausibly, that the only way to account for armchair philosophical knowledge in terms of understanding is to posit a link between understanding and assent that renders substantive disagreement impossible. But clearly this is much stronger than necessary. In general, to base the epistemology of an armchair discipline on certain cognitive capacities does not require us to assume that those capacities provide any *guarantee* that their possessor comes to know the truths to which they provide access. Consider the armchair discipline of analyzing chess problems, for example. It is quite plausible that the analysis of chess problems is based on a complex collection of cognitive capacities—one that includes, for example, grasp of the rules of chess, the ability to calculate the values of pieces, the ability to recognize abstract structural features of board positions, and the ability to calculate variants. Call this complex set of capacities *chess competence*, and let us define a chess competence-assent guarantee as follows:

CA A chess competence-assent guarantee holds for a thought t if and only if necessarily, anyone who has chess competence assent to t .

Chess players rely on their chess competence to come to know the solutions to chess problems. Yet it would be quite implausible to construe them as trying to find truths that satisfy the condition in CA. For example, for a long time chess experts generally thought that a chess endgame in which White has a king and two bishops, and Black has a king and a knight is draw. That is, despite the apparent advantage of White it was thought always to be possible for Black to find moves to prevent White from winning. For this reason, in endgames with this structure even gifted chess players playing White often would have simply accepted a draw. However, it was eventually discovered that in this endgame it is, in fact, always possible for White to checkmate Black. Hence chess competence does not guarantee assent to the truth *that in such-and-such endgames the player with two bishops can always win*. And even though the fact of the matter is conclusively settled by the rules of chess, there was no guarantee that anyone—much less everyone—who is competent in chess would ever come to know this truth. Still, there is no serious question whether the truth is one that is known purely on the basis of players' competence in chess. For the judgment to be known on the basis of chess competence, it needn't be that everyone who disagreed was incompetent.

An analogous point can be made concerning the armchair discipline of mathematics. Inquiry in mathematics also relies on a certain set of cognitive capacities, capacities that do not guarantee assent to many interesting mathematical truths. Proving Fermat's last theorem, for example, was no trivial matter, and prior to its proof some mathematicians—having tried to prove it and failed—might very well have doubted its truth. So there is no mathematical competence-assent guarantee for Fermat's theorem. It hardly follows that the theorem somehow goes beyond what can be known on the basis of mathematical competence alone, or that

Footnote 30 continued

for how it is possible to have a priori knowledge of very simple logical truths and inference rules. But Boghossian clearly does not intend to generalize this account to all armchair philosophical knowledge, as CP3 does. Nor is it plausible that armchair philosophical knowledge in general is epistemically on par with our knowledge of very simple logical truths and inference rules.

something over and above mathematical competence must be added in order to account for our knowledge of it.

These analogies suggest a picture of understanding as being (or involving) a set of cognitive capacities, just as chess and mathematical competence are. In fact, many Conceptualists do think of understanding as being (or involving) a set of cognitive capacities. According to Chalmers and Jackson, for example, a subject's grasp of a concept gives her the ability to identify the actual extension of the concept when she is given sufficient information about the actual world, and this extends to the ability to identify the concept's extension relative to arbitrary bodies of hypothetical information about ways the actual world might be.³¹ On Peacocke's view, many concepts are associated with certain informational mental states—"implicit conceptions"—and understanding such a concept gives one the ability to use the information in its corresponding implicit conception to guide one's reasoning and one's responses to sensory input.³² While these views differ in some ways about the nature of understanding, they both have the highly plausible consequence that understanding a concept is partly a matter of having an ability to identify its actual and possible instances in certain circumstances. For example, it is quite plausible that understanding the concept expressed by 'knowledge' involves being able to classify individual actual and hypothetical subjects as having, or failing to have, knowledge once one has been given certain sorts of further information about those subjects. Moreover it is plausible that understanding also involves, for example, the grasp of semantic categories and of structure, the ability to make certain logical inferences, and perhaps also the ability to engage in various kinds of inductive and abductive reasoning.

As with chess and mathematical competence, the cognitive capacities that constitute understanding need not be ideal or infallible, and it is consistent with having them that one sometimes makes mistakes: one can be misled by distracting features of a case, or by a strongly held theory that delivers the wrong verdict about it, one can be overwhelmed by the complexity of the case, one can miss subtle but important differences between it and other cases one has considered, one can make all sorts of errors of reasoning, and so on. So there is no guarantee that necessarily, everyone who understands the concept expressed by 'knowledge' is disposed to judge any particular case one way or another; much less is there any guarantee that everyone who understands it is disposed to accept any philosophically interesting truths about knowledge. Nevertheless, imperfect as they may be, one can put these sorts of cognitive capacities to work in order to acquire knowledge from the armchair, just as the chess expert and mathematician do with the capacities that constitute chess and mathematical competence.

How can the cognitive capacities that constitute understanding yield knowledge even in the absence of understanding-assent guarantees? To answer this question fully would require providing a positive Conceptualist epistemology for philosophy, something that goes far beyond the more modest defensive aims of this discussion. But it is helpful to see, at least in broad outline, how such an account might

³¹ Chalmers and Jackson (2001, pp. 323–324).

³² Peacocke (2003, 2008).

proceed.³³ Consider again the ability to classify cases that comes with understanding of the concept expressed by ‘knowledge’ and other concepts. One might systematically apply this ability to various actual and hypothetical subjects and begin to notice patterns. For example, one might notice that among all the subjects one has considered, all of them who have known *that p* have also had a true belief *that p*. This provides some justification for the hypothesis that knowledge requires true belief, and one can strengthen this justification by consider further cases of different sorts, seeking out the cases that seem most likely to provide counterexamples if there are any, trying to rule out alternative hypotheses, and so on. Such a process of inquiry can lead one to have good epistemic grounds for the belief that knowledge requires true belief, just as analogous inductive processes do in other domains. Similarly, when one applies one’s ability to Gettier subjects, one can acquire good epistemic grounds for believing that justified true belief is not sufficient for knowledge.³⁴ These processes of inquiry can provide fallible epistemic justification because the judgments we make on the basis of understanding are likely, though not guaranteed, to be true. There are different explanations that can be given for why these processes have the truth-conduciveness necessary for justification. According to one explanation, beliefs formed on the basis of understanding are at least sometimes justified because the exercise of one’s understanding-based capacities is a process that reliably tracks the instantiation of the property or properties in question across actual and possible cases.³⁵ According to another explanation, the verdicts about cases that one’s understanding-based capacities deliver contribute to determining which property the concept picks out in such a way that those verdicts tend to come out true.³⁶ No doubt there are other explanations that could be given as well. The present point is simply that cases like Williamson’s Peter and Stephen do nothing at all to show that this sort of account cannot succeed. One can make the cognitive capacities that are constitutive of understanding the epistemic basis for philosophical armchair knowledge without incurring any commitment at all to understanding-assent guarantees.

It should be emphasized that even if there are no understanding-assent guarantees, it does not follow that when one *does* come to know a thought on the basis of understanding, something *more than* understanding must have been required for knowledge. There need not be anything additional, such as a faculty of a priori insight or rational intuition, to fill the gap between understanding and assent. One’s chess competence includes an ability to calculate variants in order to identify

³³ For a more detailed positive proposal about how we can acquire knowledge via conceptual analysis, see Magdalena Balcerak Jackson (in review).

³⁴ It may be that an account along these lines will not deliver the result that understanding-based knowledge of philosophical truths is, in general, *certain* knowledge. But in our view this is a welcome result, one that is much more faithful to the actual epistemic situation of armchair philosophers. Prior to the discovery of Gettier cases, for example, philosophers had quite strong understanding-based epistemic grounds for thinking that justified true belief is sufficient for knowledge. Their grounds were epistemically on a par with the grounds we now have for thinking that knowledge is sufficient for justified true belief.

³⁵ See, for example, Jenkins (2008, 2010).

³⁶ See, for example, Peacocke (2003, 2004).

the optimal move at various points in the game. In any normal person this ability is limited and fallible: one can only calculate so-and-so many moves ahead, and even within those limits things can go wrong in a great variety of ways. But this does not mean that when things go right there must have been some special faculty in addition to one's chess competence that led one to the optimal move. The same is true for understanding: the fact that things can go wrong does not show that something in addition to understanding must be doing the work when things go right.³⁷

In our response to Williamson's argument we have been relying on an analogy between chess and mathematical competence and understanding. However, one might object to the comparison between chess competence and understanding as follows. (One could try to raise an objection to the analogy with mathematics along similar lines.) Even if chess competence is primarily a collection of cognitive capacities, there surely are *some* beliefs that a player must have in order to qualify as competent—beliefs about how the pieces move, for example, and how they are to be arranged on the board. One might argue that such secure beliefs are a necessary starting point for finding the answers to chess puzzles, even if the answers in harder and more interesting cases are not ones for which competence guarantees assent. But if Williamson is right then there are *no* beliefs that a subject must have in order to count as grasping any particular word or concept, and hence nothing that could form a secure starting point for armchair philosophical investigation.

It is certainly right that a player does not count as competent at chess unless she can follow the rules of the game when she plays. Perhaps one could insist that being able to follow the rules of chess requires having certain beliefs about how the pieces move, how they are arranged on the board, etc. and hence that chess competence constitutively involves such beliefs. But it is quite plausible that one does not count as a competent speaker of English unless one can follow certain rules as well, for example rules for constructing sentences of various sorts.³⁸ Does the ability to follow the rules in this case require that the speaker have certain beliefs? If not, then there is no constitutive connection in general between being able to follow rules and having certain beliefs. But then there is no clear reason why one could not qualify as competent in chess, and be able to employ one's competence in investigating chess problems, even in the absence of any basic beliefs for which a chess competence-assent guarantee holds. On the other hand, if being able to follow the rules of English does require having certain beliefs, then we should be very suspicious of Williamson's inference from the Peter and Stephen cases to the conclusion that there are *absolutely no* beliefs that are constitutive of understanding. For this amounts to the conclusion that one can be a competent speaker of English even without being able to follow any of the rules that govern the language. Certainly Peter and Stephen do nothing to show that such a thing is possible.³⁹

³⁷ Thanks to an anonymous referee for *Philosophical Studies* for pressing us on this point.

³⁸ An analogous claim is plausible for grasping a scheme of concepts. It is plausible that there are rules concerning, e.g. how concepts can be combined to form truth-evaluable thoughts. Concerns about whether it is appropriate to talk about "rules" in this case are tangential.

³⁹ If being able to follow the rules of English does require having certain beliefs, it may still be difficult to see how such beliefs—e.g. about the syntactic and semantic structure of the language—could help lead

Where does this leave us? As Williamson acknowledges, his discussion of Peter and Stephen relies on the sorts of intuitions that make plausible the Quinean thesis of epistemological holism, according to which any belief can be rationally given up as long as one makes appropriate adjustments to one's overall belief system.⁴⁰ Peter and Stephen are meant to illustrate that even the *prima facie* unrevisable belief that all vixens are vixens can be rejected, as long as one is prepared to modify one's other beliefs in a way that makes that rejection intelligible. What we have been arguing here is that this observation is compatible with the project of basing armchair philosophical knowledge on understanding. Understanding-assent guarantees are not needed: in general, a cognitive capacity can provide epistemic access to truths without it being such that necessarily, everyone who has the capacity assents to those truths. So even if Williamson's argument succeeds in showing that there are no understanding-assent guarantees, it falls far short of showing that nothing is "epistemically available simply on the basis of linguistic and conceptual competence."

This is a simple point, but it seems to be surprisingly easy to overlook. If we suggest, for example, that one can be justified purely on the basis of one's chess competence in judging that bringing out the knight is the best move for black in a certain sort of chess position, it would seem silly to object by pointing out that we can imagine a competent player choosing a different move. But if we suggest that the belief that chairs are furniture can be justified purely on the basis of understanding, many philosophers will want to evaluate this claim by seeing whether they can imagine a possible subject who grasps the thought without accepting it; hence they read into the suggestion a commitment to an understanding-assent guarantee for the thought *that chairs are furniture*. Why the difference? In Williamson's case, at least, we suspect that the reason is that he thinks of understanding as something profoundly unlike chess competence, and as something that is not properly construed as a cognitive capacity at all. If understanding is not a cognitive capacity, then it is not a cognitive capacity that can guarantee assent to anything. But more importantly, it is not a capacity that a thinker could, in principle, employ to try to find answers to philosophical questions, and if so then Conceptualism would certainly appear to be doomed. Hence to get to the heart of Williamson's disagreement with Conceptualism we need to address this conception of understanding directly. This is the task of the next section.

Footnote 39 continued

to answers to interesting philosophical questions. However, the same thing can be said about the beliefs (allegedly) constitutive of being able to follow the rules of chess; there is an enormous gap between simple beliefs about how the pieces move and interesting truths about complex and abstract board structures. It is worth pointing out that in both cases there are plausibly "higher-order" beliefs and inferential abilities whose presence is required in order for the relevant beliefs to constitute one's ability to follow the rules. Such beliefs and abilities would be part of one's understanding of English or chess, and as such could play a (potentially quite substantial) role in armchair investigation.

⁴⁰ p. 91.

4 Radical externalism about understanding

We can begin to get a sense of Williamson's assumptions about the nature of understanding by examining why he thinks Peter and Stephen qualify as understanding the thought *that every vixen is a vixen*. In particular, given Peter's and Stephen's deviant beliefs about universal quantification, why should we think that they properly grasp the concept expressed by (our use of) 'every'? Williamson cites several intuitive considerations: that they learned English in the normal way, and at least initially used 'every' in the same way as other English speakers; that having become convinced of unorthodox theories should not rob them of the understanding they previously had; that their deviant beliefs make little practical difference to their day-to-day communicative interactions with others; and that both emphatically insist that they intend to be using 'every' the same way as other English speakers, and would be prepared to revise their usage if they could be convinced that they were wrong. As Williamson summarizes these considerations: "They joined the club of 'every'-users; since they haven't resigned or been expelled, they are still members."⁴¹

The picture suggested by these observations, and by this remark in particular, clearly brings to mind Burge-style social externalism, and Williamson acknowledges as much.⁴² In fact, he thinks that standard forms of social externalism don't go far enough, at least not when they still allow a distinction between experts with full understanding and ordinary members of the community that defer to the experts:

...experts themselves can make deviant applications of words as a result of theoretical errors and still count as fully understanding their words. ...Their assignments of meaning to those words are not parasitic on the assignments that more privileged individuals make. Rather, each individual uses words as words of a public language; their meanings are constitutively determined not individually but socially, through the spectrum of linguistic activity across the community as a whole. The social determination of meaning requires nothing like an exact match in use between different individuals; it requires only enough connection in use between them to form a social practice. Full participation in that practice constitutes full understanding. That is why there is no litmus test for understanding.⁴³

To say that "there is no litmus test for understanding" is to say that there are no particular beliefs that are required for membership in the linguistic community:

There is, of course, a distinction between understanding a word and not understanding it. One can lack understanding of a word through lack of causal interaction with the social practice of using a word, or through interaction too superficial to permit sufficiently fluent engagement in the practice. But

⁴¹ pp. 90–91.

⁴² p. 91.

⁴³ p. 98.

sufficiently fluent engagement in the practice can take many forms, which have no single core of agreement.⁴⁴

The conception suggested here is radically externalist. Understanding a language—and grasping the concepts and thoughts expressible in the language—requires merely that one have the right sorts of causal interactions with other members of the linguistic community.

If this conception of understanding is correct, then it is extremely plausible that there are no understanding-assent guarantees. Peter and Stephen can count as grasping the thought that every vixen is a vixen, despite their deviant beliefs, because they have “sufficiently fluent engagement” in the practices of using ‘vixen’, ‘every’ and so on. For any putative understanding-assent guarantee we should, in principle, always be able to find a hypothetical subject who violates the guarantee but who has sufficient causal interaction with the relevant linguistic practice to count as understanding the relevant sentences. Thus Williamson’s conception of understanding implies that we should expect there to be no understanding-assent guarantees that survive scrutiny.

But more importantly, Williamson’s conception implies that understanding is just not the sort of thing that can plausibly form the epistemic basis for armchair philosophical knowledge. Understanding is not any sort of cognitive capacity, on this view, but rather a matter of the relations one bears to others—a matter of whether one fits appropriately into a community of language or concept users. Fitting in merely requires sufficient causal interaction with the relevant social practices, and need not equip one with any intellectual tools whatsoever, or at least none that would be of any use from the armchair.⁴⁵

This radical externalist conception of understanding is clearly at odds with Conceptualism, and if it is correct then Williamson is right that we should look elsewhere for a justification of armchair philosophy. But why should we think that this is the right conception of understanding? Williamson claims that his argument that Peter and Stephen understand the thought *that every vixen is a vixen* appeals “to features of our ascription of beliefs that make semantic externalism plausible.”⁴⁶ But frankly, we doubt that most philosophers have a very strong conviction that Peter and Stephen do understand, and these cases do not, by themselves, lend much support to such a radical theory of understanding. And it certainly *is* a radical theory. For one thing, it seems quite plausible that understanding an expression or concept brings with it *some* kind of ability to make judgments about actual and possible cases, and there would appear to be no room for this suggestion on Williamson’s view. For another, one’s understanding of language at least usually contributes to knowledge of what others have said, and then sometimes on to further knowledge on the basis of what they have said. It is hard to see how understanding

⁴⁴ p. 126.

⁴⁵ In conversation, Williamson acknowledges that fitting into a community of language or concept users may turn out to require that a subject satisfy some sort of “higher-order” rationality constraints. There is no discussion in *The Philosophy of Philosophy* of why we should expect the contribution to armchair philosophical knowledge made by such constraints to be negligible. (See also n. 39 above.)

⁴⁶ p. 91.

can play this role if it is merely a matter of causal interactions with other speakers.⁴⁷ Nor is it easy to see how understanding in this sense can help capture the evident fact that assertions are intentional actions, done deliberately and guided partly by our understanding of what our sentences mean.⁴⁸ Moreover, understanding in Williamson's sense seems to have nothing to contribute to an explanation of the familiar and striking phenomenology of understanding. Why should coming to stand in certain causal relations with French speakers, for example, make it the case that one begins to hear utterances of French sentences differently?⁴⁹ How can participation in the social practices of *p*-thinkers and *q*-thinkers explain the fact that there is something it is like to grasp the thought *that p* that is distinct from what it is like to grasp *that q*?⁵⁰ Adopting Williamson's conception of understanding would seem to require giving up on the idea that understanding has a significant explanatory role to play in any of these areas.

A full attempt to adjudicate the dispute about the nature of understanding using these sorts of considerations is beyond our current scope. But it is not hard to see that a Conceptualist conception that sees understanding as a cognitive capacity—or, better, as a complex set of such capacities, as suggested in the previous section—at least looks more promising than Williamson's radical externalism in these respects. Consider, by analogy, the fact that the cognitive capacities that constitute chess competence can make it rational to make certain moves in a chess game, and to interpret the moves of others in certain ways. Consider also the fact developing one's chess capacities has quite noticeable effects on the character of one's visual experiences of arrangements of pieces on the board.

In fact, we think the analogy with chess should be taken even further. Chess competence comes in degrees: some people are better than others at chess, because people have the relevant cognitive capacities to greater or lesser degrees, are able to apply them more or less effectively, and are able to integrate them more or less well. Hence it is actually quite artificial to speak of chess competence as something that one either has or lacks, as we did when formulating the notion of a chess competence-assert guarantee in CA above. Should we understand CA as quantifying only over subjects who rise above a certain threshold level of competence? How do we choose the level? Or do we have different chess competence-assert guarantees for different levels of competence? The very idea of a chess competence-assert guarantee begins to look unmotivated. Which truths one accepts on the basis of one's chess competence depends on what kind of competence one has and on how one uses it. One could try to motivate a certain non-arbitrary minimal level of chess competence. For example, as we noted above, there is a clear difference between those who can follow the rules of chess and those who can't. Minimal competence in this sense is what allows one to play chess, in a very basic sense, but it doesn't explain much else. Minimal competence is of virtually no interest when we want to investigate how people solve interesting chess problems.

⁴⁷ See Fricker (2003).

⁴⁸ See Heck (2006).

⁴⁹ See Strawson (1994, Chap. 1) and Pettit (2001, Chap. 2).

⁵⁰ See Pitt (2004).

By analogy, it is quite plausible that the cognitive capacities that are constitutive of understanding also come in degrees, and hence that it is artificial to speak of understanding as something that one either has or lacks. As with chess, we can try to motivate a certain non-arbitrary minimal level of competence. Williamson argues that understanding in his sense is what matters for “smooth and fruitful interaction with other members of the community,” because it is,

...an appropriate response to an important constraint on a theory of linguistic meanings: that there is little point in talking about them unless they can be shared across significant differences in belief, between different individuals at the same time or the same individual at different times. They can survive factual learning and factual disagreement.⁵¹

We hasten to point out that nothing about a conception of understanding in terms of cognitive capacities entails that communication across differences in belief is impossible, or even problematic. But perhaps Williamson’s radical externalist conception is the best way to capture understanding in this minimal sense. If so then Williamson’s argument is, in effect, that minimal competence in this sense has no substantial explanatory role to play in the methodology of philosophy. It probably doesn’t have a substantial explanatory role to play in very many other areas either, though, and the Conceptualist is certainly under no obligation to accept it as the only notion of understanding that is philosophically significant.

Where does this leave us with Peter and Stephen—do they understand the concept expressed by ‘every’ or not? Given that we see understanding as coming in degrees, it should come as no surprise that we think there is no simple ‘yes’ or ‘no’ answer here. Peter and Stephen understand in Williamson’s minimal sense (as they are designed to) and hence perhaps they at least have an “entry-level” understanding. However, they may not understand the concept as well as some others do. For example, many people who understand the concept recognize that whenever, say, the thought *that every dog barks* is true then the thought *that every brown dog barks* must be true as well, and similarly for many other inferences of this form. But Peter must regard this inference pattern as invalid, because it is consistent with every dog barking that there aren’t any brown dogs, and hence that *that every brown dog barks* is false. Stephen must reject it as well, because there may be borderline cases of brown-doghood even if there are no borderline cases of doghood. For similar reasons Stephen must also reject the inference from the thought *that every dog barks loudly* to the thought *that every dog barks*, something that many subjects who grasp the concept expressed by ‘every’ recognize as valid. These deviances are much more widespread, and don’t depend on a small cluster of unusual beliefs about vixens, and hence they are at least *prima facie* reason to regard Peter and Stephen as having a poorer grasp of the concept than many ordinary subjects.

Nevertheless, perhaps we can imagine that Peter and Stephen recognize the intuitive pull of these inference patterns, but take their theoretical beliefs to give them good reason to reject them in the final analysis. Perhaps they grasp the concept as well as one can ever reasonably hope, and yet still fail to assent to the thought

⁵¹ p. 97.

that every vixen is a vixen. We have argued that this is a possibility that Conceptualists can and should live with. Even “expert understanders” can be seduced by misleading theories, as can mathematicians and chess experts.⁵²

5 Conclusion

We conclude by considering a question: on the view we have defended here, are there in fact any conceptual truths? The short answer to this question is ‘yes’. The thought *that every vixen is a vixen* is a conceptual truth, as are the thoughts *that every bachelor is male* and *that chairs are furniture*. We are justified in accepting these thoughts as true, and our justification for doing so rests on our understanding of the concepts involved and the ways they are combined. Since they can be justified in this way, in our view it is fine to label them all as conceptual truths. Slightly more interestingly, we think the thought *that knowledge is not justified true belief* is a conceptual truth, as is the thought *that a person can survive the loss of her brain*. These cases are somewhat more controversial, although as Williamson shows even the apparently trivial ones can become controversial in the right circumstances.

However, in classifying these thoughts as conceptual truths it is absolutely crucial to avoid falling into certain confusions. To say that they are conceptual truths is not at all to say that they are merely metalinguistic or meta-conceptual truths. On the account we sketched in Sect. 3, understanding the concept expressed by ‘knowledge’ brings with it a capacity to classify particular cases as cases of knowledge or not. The reflective exercise of this capacity allows one to make discoveries about knowledge itself, not just about the word ‘knowledge’ or the concept expressed by it. Nor does the account make use of any assumption that the truths discovered are somehow mere “truths of language” that do not make any substantial claims about the world. Conceptual truths in our sense are truths that are in principle knowable purely on the basis of understanding. They are not distinguished by any special, distinctive subject-matter at all. They are truths about knowledge, free will, personal identity or anything else philosophers happen to have an interest in investigating. They are distinguished only by the sort of epistemic access we have to them. Recognizing the viability of such an attitude towards conceptual truths is crucial for developing a systematic epistemology of armchair philosophy based on understanding.

Acknowledgments Earlier versions of this material were given at the Thursday seminar of the Philosophy Program at the Research School of the Social Sciences, Australian National University and at a workshop at the University of Cologne, Germany. We are grateful to audiences at both events for much helpful discussion. Thanks especially to David Chalmers, Thomas Grundmann, Andrew McGonigal, Daniel Stoljar and Timothy Williamson. The research for this essay was supported by a grant from the Deutsche Forschungsgemeinschaft.

⁵² A further difficult question, and one on which we take no stand here, is whether a subject with an *ideal* understanding of the concept (in some hard-to-specify sense) can fail to recognize the truth of the thought *that every vixen is a vixen*.

References

- Ayer, A. J. (1946). *Language, truth & logic* (2nd ed.). New York: Dover Publications, Inc.
- Ayer, A. J. (1954). Freedom and necessity (Reprinted in *Philosophical Essays*, by A. J. Ayer).
- Balcerak Jackson, M. (2011). Manuscript. "Conceptual analysis and epistemic progress" (in review).
- Boghossian, P. (1994). Analyticity and conceptual truth. *Philosophical Issues: Truth and Rationality*, 5, 117–131.
- Boghossian, P. (1996). Analyticity reconsidered. *Noûs*, 30, 360–391.
- Boghossian, P. (2003). Blind reasoning. *Aristotelian Society Supplementary*, 77, 225–248.
- Chalmers, D. J., & Jackson, F. (2001). Conceptual analysis and reductive explanation. *Philosophical Review*, 110, 315–360.
- Fricker, E. (2003). Understanding and knowledge of what is said. In A. Barber (Ed.), *Epistemology of language*. Oxford: Oxford University Press.
- Goldman, A. (2007). Philosophical intuitions: Their target, their source, and their epistemic status. *Grazer Philosophische Studien*, 74, 1–26.
- Goodman, N. (1954). *Fact, fiction and forecast*. London: Athlone Press.
- Heck, R. (2006). Reason and language. In G. MacDonald & C. MacDonald (Eds.), *McDowell and his critics*. Oxford: Blackwell Publishing.
- Hempel, C. (1945). Studies in the logic of confirmation I. *Mind*, 54, 1–26.
- Jackson, F. (1998). *From metaphysics to ethics*. Oxford: Oxford University Press.
- Jenkins, C. S. (2008). *Grounding concepts: An empirical basis for arithmetical knowledge*. Oxford: Oxford University Press.
- Jenkins, C. S. (2010). Concepts, experience and modal knowledge. *Philosophical Perspectives*, 24, 255–279.
- Peacocke, C. (2003). Implicit conceptions, understanding, and rationality. In M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler Burge*. Cambridge: MIT Press.
- Peacocke, C. (2004). *The realm of reason*. Oxford: Oxford University Press.
- Peacocke, C. (2008). *Truly understood*. Oxford: Oxford University Press.
- Pettit, D. R. (2001). *Understanding language*. PhD Dissertation, Massachusetts Institute of Technology, Cambridge.
- Pitt, D. (2004). The phenomenology of cognition, or, what is it like to think that P? *Philosophy and Phenomenological Research*, 69, 1–36.
- Russell, G. (2008). *Truth in virtue of meaning*. Oxford: Oxford University Press.
- Strawson, G. (1994). *Mental reality*. Cambridge: MIT Press.
- Williamson, T. (2003). Understanding and inference. *Aristotelian Society Supplementary*, 77, 249–293.
- Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell Publishing.